

AD736213

Report No. 2309

January 1972

INTERFACE MESSAGE PROCESSORS FOR
THE ARPA COMPUTER NETWORK


QTR No. 11 - AD 731687

QUARTERLY TECHNICAL REPORT NO. 12
1 October 1971 to 31 December 1971

Principal Investigator: Mr. Frank E. Heart
Telephone (617) 491-1850, Ext. 470

Sponsored by
Advanced Research Projects Agency
ARPA Order No. 1260


Contract No. DAHC15-69-C-0179
Effective Date: 2 January 1969
Expiration Date: 31 December 1972
Contract Amount: \$5,454,787



Title of Work: IMP

Submitted to:

Director
Advanced Research Projects Agency
Arlington, Virginia 22209



Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
Springfield, Va. 22151

19

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Bolt Beranek and Newman Inc. 50 Moulton Street Cambridge, Mass. 02138		2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	
		2b. GROUP	
3. REPORT TITLE QUARTERLY TECHNICAL REPORT NO. 12			
4. DESCRIPTIVE NOTES (Type of report and inclusive dates)			
5. AUTHOR(S) (First name, middle initial, last name) Bolt Beranek and Newman Inc.			
6. REPORT DATE January 1972		7a. TOTAL NO. OF PAGES 14	7b. NO. OF REFS
8a. CONTRACT OR GRANT NO. DAHC15-69-C-0179		9a. ORIGINATOR'S REPORT NUMBER(S) BBN Report No. 2309	
b. PROJECT NO. 1260		9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)	
c.			
d.			
10. DISTRIBUTION STATEMENT			
11. SUPPLEMENTARY NOTES		12. SPONSORING MILITARY ACTIVITY Advanced Research Projects Agency Arlington, Virginia 22209	
13. ABSTRACT The basic function of the ARPA computer network is to allow large existing computers (Hosts), with different system configurations, to communicate with each other. Each Host is connected to an Interface Message Processor (IMP), which transmits messages from its Host(s) to other Hosts and accepts messages for its Host(s) from other Hosts. There is frequently no direct communication circuit between two Hosts that wish to communicate; in these cases intermediate IMPs act as message switchers. The message switching is performed as a store and forward operation. The IMPs regularly exchange information which: allows each IMP to adapt its message routing to the conditions of its local section of the network; reports network performance and malfunctions to a Network Control Center; permits message tracing so that network operation can be studied comprehensively; allows network reconfiguration without reprogramming each IMP. The Terminal IMP (TIP), which consists of an IMP and a Multi-Line Controller (MLC), extends the network concepts by permitting the direct attachment (without an intervening Host) of up to 64 dissimilar terminal devices to the network. The Terminal IMP program provides many aspects of the Host protocols in order to allow effective communication between a terminal user and a Host process.			

14 KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Computers and Communication						
Store and Forward Communication						
ARPA Computer Network						
Interface Message Processor						
IMP						
Terminal IMP						
TIP						
Honeywell DDP-516						
Honeywell DDP-316						
Multi-Line Controller						
MLC						
Network Control Center						
Host Protocol						
High Speed Modular IMP						
HSMIMP						
Reliable Transmission Package						

INTERFACE MESSAGE PROCESSORS FOR
THE ARPA COMPUTER NETWORK

QUARTERLY TECHNICAL REPORT NO. 12
1 October 1971 to 31 December 1971

Submitted to:

Advanced Research Projects Agency
Washington, D.C. 22209
Attn: Dr. L.G. Roberts

This research was supported by the Advanced Research Projects Agency of the Department of Defense under Contract No. DAHC-15-69-C-0179.

TABLE OF CONTENTS

	Page No.
1. OVERVIEW	1
2. TERMINAL IMP	6
3. "VERY DISTANT HOST" INTERFACE	9

1. OVERVIEW

This Quarterly Technical Report, Number 12, describes aspects of our work on the ARPA Computer Network during the last quarter of 1971.

During this quarter the first Model 316 IMP was installed in the field and two additional Terminal IMPs were delivered. The Model 316 IMP, which was installed at ETAC (Environmental Technical Applications Center, Washington, D.C.), will be replaced by a Terminal IMP during the first quarter of 1972. The performance of the installed Terminal IMPs is discussed in Section 2.

During the fourth quarter several documents were completed and distributed. They include:

- BBN Report No. 1822, *Specifications for the Interconnection of a Host and an IMP*; a major revision of this document was released in October.
- BBN Report No. 2183, *User's Guide to the Terminal IMP*; a major revision of this document was released in December.
- *The Terminal IMP for the ARPA Computer Network* (a paper submitted to the 1972 Spring Joint Computer Conference); this paper describes the Terminal IMP in some detail and also discusses recent developments in the IMP subnetwork. This paper was distributed to the Network community in December.

Almost since the installation of the first IMPs, sites have been dissatisfied with the 2000 foot cable limitation inherent in the design of the "distant Host" interface. Longer cable lengths

appear feasible only if error detection and recovery procedures are used. During this quarter we have designed a new type of Host interface which incorporates this type of line control. This "very distant Host" interface is described in Section 3.

A major effort during this quarter was the design and partial coding of improved methods of flow control and lockup prevention in the IMP subnetwork. The current subnetwork implementation is subject to various kinds of congestion which can degrade performance under heavy load. Although this has not been an operational problem with current Network traffic, methods of preventing these situations from arising have been the subject of intensive study. We believe it will be possible to implement congestion control mechanisms in the Network during early 1972.

A number of current or prospective Terminal IMP sites, as well as some of the Network's large Host "service sites", would like to see Terminal IMPs provide for the connection of "remote batch" terminals to ports on the Multi-Line Controller (MLC). Most commercially available remote batch terminals transmit data blocks of several characters with only block boundaries marked to obtain high bandwidth on their communication channels. The MLC hardware, however, requires the beginning and end of each character to be marked (with "start" and "stop" bits) so that common types of remote batch terminals cannot be directly connected to the MLC. Further, in conjunction with the data-block format, remote batch terminals usually employ block checksum, acknowledgment, and re-transmission schemes to detect communication line errors. Not only is there limited TIP core storage and program bandwidth for implementing any of these schemes but, in addition, there are several different methods in common use.

During this quarter, however, we have arrived at a tentative solution to the problem of connecting remote batch terminals to the MLC based on a new type of modem which *by itself* completely handles line-error detection and correction by means of buffering and retransmission logic internal to the modem. This modem could be used with any one of the several types of remote batch terminals which are built around mini-computers of various kinds, and therefore could be programmed to provide the bit configuration required by the MLC. During the early part of 1972 we expect to obtain the use of one of these terminals and, after the appropriate reprogramming is performed, experiment with the connection of it to the MLC via an error-correcting modem.

During the fourth quarter work continued on the attachment of magnetic tape drives to Terminal IMPs. We have now received two drives from Honeywell, along with the associated controllers, cabinets, and additional core memory modules. By the end of the quarter programming was close to completion and some testing was underway in the BBN test cell. Field installation of the first Terminal IMP equipped with magnetic tape is scheduled for the first quarter of 1972.

Fabrication of two Univac 418 special Host interfaces (see BBN Report No. 2270, *Quarterly Technical Report No. 11*) was completed during this quarter and the interfaces have been tested as completely as possible in a "stand alone" mode.

Each IMP in the Network was taken down for hardware retrofitting during the month of November. Corrections were made in several details of the IMP design which included:

- A minor electrical noise problem in the transmit side of the (IMP-to-IMP) modem which occasionally caused the modem to appear hung.
- A data-clocking problem on the IMP output bus which caused occasional IMP-to-Host communication failures with some Hosts.
- The electrical characteristics of the signals on some interrupt lines were changed to increase noise rejection on the lines.
- The power-failure autorestart feature was added to the Terminal IMPs already in the field.

Early in 1971 we began to consider methods for developing a new version of the IMP which would be significantly faster than the present IMP and more modular in nature. Its higher speed would permit it to take advantage of the high speed (i.e., megabit and above) communication lines which the common carriers are currently developing. More generally, this new type of IMP would be able to service high-bandwidth requirements wherever and however they occur in the Network. Further, the goal of modularity is intended to lead to increased versatility. This new type of IMP is referred to as the High Speed Modular IMP, or HSMIMP.

Our current thinking about the design of the HSMIMP is that a multiprocessor configuration of minicomputers, each with some private memory as well as access to a common memory, is most sensible. I/O multiplexors and CPUs will access the common memory through a high-speed crossbar switch. We expect to adapt some commercially-available mini-computer, to design the switch and the I/O multiplexor interfaces, and to use commercially-available memories.

During this quarter we began to actively study several commercially-available system components which might be appropriate for use in the HSMIMP. In addition, we have examined the architecture of a number of proposed or existing minicomputer-multi-processor systems to determine their suitability for construction of the HSMIMP. We expect the pace of work on the HSMIMP to pick up sharply in the first quarter of 1972.

Finally, during the fourth quarter we continued our "login" survey of the Network's "service" Hosts from the BBN prototype Terminal IMP. Since the middle of September we have been able to successfully log in to nine Hosts. Once a login was successfully completed to a given Host, that Host was surveyed once a day on most weekdays, and a "percentage availability" figure calculated. For these nine Hosts this percentage ranged from 22% to 72%.

2. TERMINAL IMP

By the end of the fourth quarter three Terminal IMPs had been operating in the field for at least one month, at MITRE, NASA/Ames, and NBS (National Bureau of Standards, Washington, D.C.). The first two of these, in fact, were installed during the third quarter. Version 5 of the TIP program (the current version) was released to the field during the second week in December. Since there has now been sufficient opportunity for significant Terminal IMP use it is appropriate to discuss initial TIP user experience in this report. Accordingly we have solicited comments from the three sites and summarize them here.

Generally speaking, all three sites seem to be relatively happy with the performance of the Terminal IMP, although NBS is less happy than the other two sites. All sites experienced several software difficulties with earlier versions of the TIP program, but felt that Version 5 has eliminated or significantly reduced these problems.

On the other hand, users were rather unhappy with two aspects of overall Network performance. First, all three sites experienced significant difficulties with service Host availability (MITRE had less complaints in this regard than the other two sites). Problems have included frequent Host crashes, loss of stored data, small number of Hosts even *attempting* to provide service on a regular basis, and, in the case of NBS, inability to communicate with some Hosts because those Hosts must modify internal tables in order to "recognize the existence" of NBS and have not done so. Second, both MITRE and NBS are in the center of a long Network loop, so that if an IMP or communication circuit on each side of one of these sites goes down, that site is isolated from the rest of the Network. NBS in particular complains of this partitioning occurring several times.

MITRE users mentioned a particular problem encountered when trying to use Multics which, when generating large volumes of output (e.g., one paragraph), tries to send it all at once. Under these conditions, the TIP has insufficient buffer capacity to store all the data which Multics wishes to send. Multics sends as much as the TIP will accept and then apparently puts the sending process to sleep. At times of heavy load on Multics it takes up to several minutes to awaken the sending process, thus causing intolerable delays in the middle of output. As a temporary expedient they gain access to Multics via either the BBN TENEX system or the MIT Dynamic Modeling system, each of which has sufficient buffering and rapid enough response to smooth out the input/output flow.

NBS users commented that they have tried to use the UCLA 360/91 several times, but have been unsuccessful thus far. The TIP is currently unable to interact directly with the Remote Job Service offered by the 360/91 (this problem is being attacked in various ways by both UCLA and BBN). Other server Hosts are also generally unable to interact with the 360/91, although the UCLA Sigma-7 can sometimes be used as an intermediary between Terminal IMPs and the 360/91. NBS, however, has been unable to establish communication with the Remote Job Service via the Sigma-7.

Those sites which commented on the MLC (Multi-Line Controller) were satisfied with its overall operation, and with the amount of information which is displayed on the control panel. One complaint voiced by MITRE was that the Line Interface Units conform to the EIA RS-232-C specification in guaranteeing to drive data lines only to a maximum of fifty feet. It was felt that this limitation unreasonably restricted the placement of terminals at their site.

It was suggested by some sites that the TIP control panel should be equipped with a lockable cover so that unauthorized and/or inadvertent changes could not be made to the switch settings.


Finally, the MITRE administrative personnel requested that thought be given to methods of keeping track of who is using (and has used) the Terminal IMP. They are currently faced with the problem of identifying terminal users who dial in to modems equipped with the automatic-answer feature. Although this is not yet a critical problem for them, future accounting policies may make it desirable to monitor and account for such use.

In summary, Terminal IMP users appear to be happy with the design and current operation of the TIP system, and less satisfied with the availability of Network resources. Several minor suggestions for improvement in the user interface and requests for new capabilities were made. One or two larger problems, although not yet critical, are likely to require study in the near future.

3. "VERY DISTANT HOST" INTERFACE

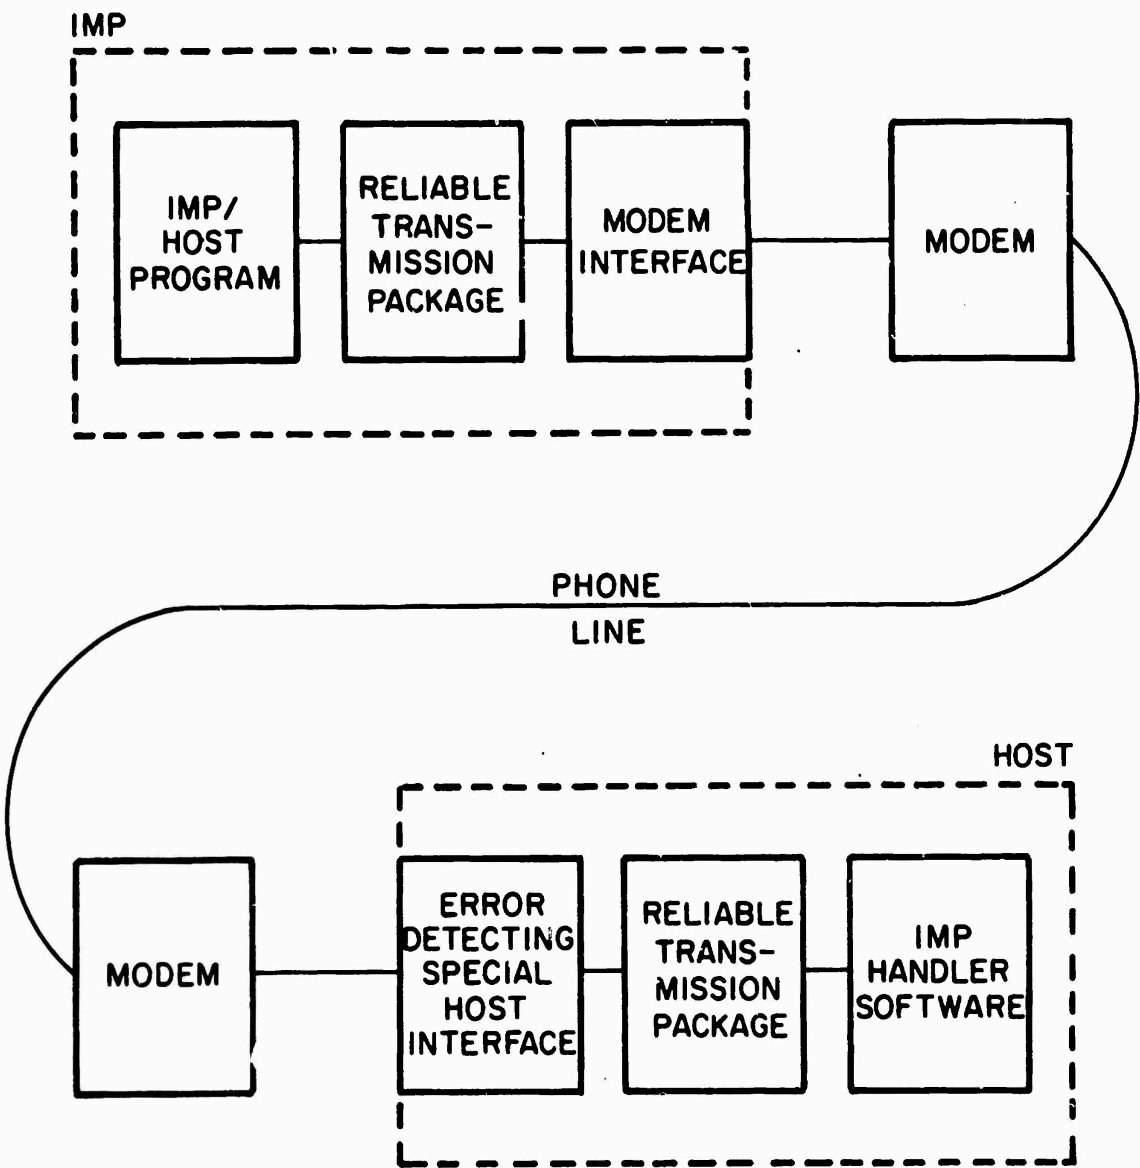
During the fourth quarter we studied the issue of connecting a Host and an IMP over a distance greater than 2000 feet, the maximum distance over which the distant Host interface can guarantee error-free transmission. We decided that such a connection was feasible with a relatively small change to the IMP, if made in the manner discussed in the following paragraphs. We call such a connection a *Very Distant Host* connection.

Currently, connection between an IMP and any of its Hosts takes place as illustrated below:



The Host (or Distant Host) Interface and the Special Host Interface communicate according to the hardware specification set down in BBN Report No. 1822, *Specifications for the Interconnection of a Host and an IMP*, and the IMP/Host program (in the IMP) and the IMP Handler software (in the Host) communicate using the software protocol described in the same document.

To minimize the disturbance to existing programs and specifications in both the IMPs and the Hosts, the Very Distant Host connections will be implemented by adding a new level of protocol (which can be programmed in self-contained front end packages) and using the IMP's standard modem interface as shown below:




At the IMP end of the connection, the Host interface is replaced by a modem interface and a modem, and a software package which provides reliable packet transmission is added between the IMP/Host program and the hardware interface. At the Host end of the connection, a modem is added, along with some sort of hardware device which provides an interface between the modem and the Host. Also, between the hardware interface and the IMP Handler software, a software package which provides reliable packet transmission is added. As before, the IMP/Host program (in the IMP) and the IMP Handler software (in the Host) communicate according to the software specifications in BBN Report No. 1822 except that the leader will be separated from the first packet of the message. The new Reliable Transmission Packages in the IMP and the Host communicate as outlined below; the modem interface in the IMP and the Error Detecting Special Host Interface communicate using the line protocol currently used between IMPs (as described in BBN Report No. 1763, *Initial Design for Interface Message Processors for the ARPA Computer Network*, pages 37 through 41).

The only new development that is required at the IMP end is the Reliable Transmission Package. At the Host end, both the Reliable Transmission Package and a new piece of hardware, namely the Error Detecting Special Host Interface, require development. We will discuss the new piece of Host hardware first.

We see three reasonable ways for the Host to build the necessary hardware.

1. Build the equivalent of the IMP modem interface, which will provide an interface between the modem and the Host.

2. Adapt the Special Host Interface so it talks to the modem instead of to an IMP, e.g.,



3. Place a mini-computer between the Host and the modem (and program the Reliable Transmission Package in the mini-computer).

Any of these methods is feasible; which one is chosen will depend on what is comfortable at the Host site.

There are several ramifications of the scheme just described which we should point out:

There is no ready line on a modem interface, and there is no need for one. Periodic "no-ops" sent between the Reliable Transmission Packages can be used to inform both the Host and the IMP of the other's "aliveness."

The IMP will require that transmission be in packets which are multiples of 16 bits, up to a maximum of 1008 bits. Thus, unlike a normal Host, a Very Distant Host must be aware of packets. Also, a word mismatch problem may exist between IMPs and Very Distant Hosts. For example, a PDP-10 would have to send out four of its 36-bit words before it hit a 16-bit boundary. Packets between an IMP and a Very Distant Host will include a count of 16-bit words to facilitate finding packet boundaries.

The current IMP/Host interface imposes no timing constraints on the Host or the IMP, since it is based on a per-bit handshake. The new interface will demand that both Host and IMP handle data at a rate dictated by a phone line modem.

We will now briefly describe the algorithm to be used by the Reliable Transmission Packages.

The Reliable Transmission Packages (RTP) in the Host and IMP are functionally equivalent. Both send and receive packets which are multiples of 16 bits in length. Appended to the front of each packet is 16 bits of control information including a count giving the length (in 16-bit words) of the data in the packet, a bit which when set indicates the last packet of a message, an odd/even bit which is used to detect duplicate packet transmissions, a one-bit "channel number", and two acknowledge bits — one for channel zero and one for channel one. Further, for efficiency, the RTPs must be able to handle two packets going in each direction (transmit and receive) simultaneously.

At any time each of the two packets going in one direction is associated with one of the two "channels" mentioned above. For each transmit channel a used/unused bit and an odd/even bit are kept (both initialized to zero). The used/unused bit indicates whether there is currently a packet associated with the channel. For each receive channel, an odd/even bit is kept (initialized to one). The transmit portion of the RTP cycles through its used channels (those with packets associated with them), transmitting the packets along with the channel number and the associated odd/even bit. At the receive side of the RTP, if the odd/even bit of the received packet does not match the odd/even bit associated with the appropriate receive channel, the receive odd/even bit

is complemented. Otherwise the packet is a duplicate and is discarded. Acknowledgments of all packets correctly received at the receive side of the RTP, whether duplicate or not, are sent to the transmit side of the RTP at the other end of the telephone line. This is done by copying both of the receive odd/even bits into the position reserved for the two acknowledge bits in the control information of *every* packet transmitted. In the absence of other traffic, the acknowledges are returned in 16 bit "null packets." These have a word count of zero and only the two acknowledge bits contain relevant information. When the transmit portion of the RTP receives a packet, it compares (bit by bit) the two acknowledge bits against the two transmit odd/even bits. For each match found, the corresponding channel is marked unused and the corresponding packet is discarded, and the odd/even bit is complemented. The transmit portion of the RTP must fill its channels in sequence (one to channel zero, one to channel one, one to channel zero, ...), waiting if necessary for any outstanding acknowledgments. The receive portion must pass on correctly received packets in sequence, waiting for the retransmission of any missed packet.

The transmit portion of the RTP declares the line dead if a packet has not been acknowledged within t seconds. In the absence of traffic the transmit portion must generate a discardable packet (a null packet will do) every r seconds. Once a line is declared dead, no acknowledges should be returned for a period of $2*(t + r)$ seconds in order to be sure the far side also learns the line is dead and reinitializes. Five seconds is a reasonable value of t and $2\frac{1}{2}$ seconds is a reasonable value of r . The odd/even bits and channel sequence must be initialized at start-up and reinitialized after every dead sequence.

Construction of the reliable transmission package for the IMP was begun this quarter.